arXiv:1702.05224v1 [cs.DM] 17 Feb 2017

Continuous Relaxations for the Traveling Salesman Problem

Tuhin Sahai[1], Stefan Klus[2], and Michael Dellnitz[3]

[1]United Technologies Research Center, 2855 Telegraph Ave, Berkeley, CA, USA
[2]Freie Universität Berlin, Arnimallee 9, 14195 Berlin, Germany
[3]Paderborn University, Warburger Straße 100, 33098 Paderborn, Germany

Abstract

In this work, we construct heuristic approaches for the traveling salesman problem (TSP) based on embedding the discrete optimization problem into continuous spaces. We explore multiple embedding techniques – namely, the construction of dynamical flows on the manifold of orthogonal matrices and associated Procrustes approximations of the TSP cost function. In particular, we find that the Procrustes approximation provides a competitive biasing approach for the Lin–Kernighan heuristic. The Lin–Kernighan heuristic is typically based on the computation of edges that have a "high probability" of being in the shortest tour, thereby effectively pruning the search space for the k-opt moves. This computation is traditionally based on generating 1-trees for a modified distance matrix obtained by a subgradient optimization method. Our novel approach, instead, relies on the solution of a two-sided orthogonal Procrustes problem, a natural relaxation of the combinatorial optimization problem to the manifold of orthogonal matrices and the subsequent use of this solution to bias the k-opt moves in the Lin–Kernighan heuristic. Although the initial cost of computing these edges using the Procrustes solution is higher than the 1-tree approach, we find that the Procrustes solution, when coupled with a homotopy computation, contains valuable information regarding the optimal edges. We explore the Procrustes based approach on several TSP instances and find that our approach often requires fewer k-opt moves.

1 Introduction

The traveling salesman problem (TSP) is an iconic NP-hard problem that has received decades of interest [1]. This combinatorial optimization problem arises in a wide variety of applications related to genome map construction [2], telescope management [3, 4], and drilling circuit boards [5]. The TSP also naturally arises in applications related to target tracking [6], vehicle routing [7], and communication networks [8] to name a few. Recently, a history dependent TSP was used to construct efficient techniques for learning the structure of Bayesian networks [9]. For further information about applications related to the TSP, we refer the reader to [1].

In its basic form, the statement of the TSP is exceedingly simple. The task is to find the shortest Hamiltonian circuit through a list of cities, given their pairwise distances. Despite

its simplistic appearance, the underlying problem is NP-hard [10]. Several heuristics have been developed over the years to solve the problem including ant colony optimization [11], cutting plane methods [12, 13], Christofides heuristic algorithm [14], and the Lin–Kernighan heuristic [15].

In this work, we concentrate on constructing novel approximations to the TSP that are based on continuous relaxations or embeddings. In the first part, we construct a dynamical systems approach for computing solutions of the TSP. This flow on the manifold of orthogonal matrices converges to a permutation matrix that minimizes the tour length. Although the approach is interesting and elegant, the flow often converges to local minima. For TSP instances with more than 50 cities, these minima are not competitive when compared to state-of-the-art heuristics [1, 12, 13, 15]. However, inspired by this continuous relaxation, we compute the solution to a two-sided orthogonal Procrustes problem [16] that relaxes the TSP to the manifold of orthogonal matrices. We find that this Procrustes approach can be combined with the Lin–Kernighan heuristic [15] to construct a state-of-the-art heuristic for computing solutions of the TSP. The Lin–Kernighan heuristic is an extremely popular approach for the TSP and has been credited with finding the best known solutions for several large instances [17, 18]. The approach has been particularly successful in finding the best known solutions for several asymmetric TSPs [17]. We provide a detailed description of the Lin–Kernighan heuristic in Section 4. A highly successful software implementation of the approach is Helsgaun's software package LKH [17] and our approach is integrated with this package. This implementation uses minimum spanning trees [19, 20] to pre-compute candidate edges that are likely to be a part of the solution. This biasing methodology is found to provide significant improvement over baseline LKH software performance [17]. In our approach, the Procrustes solution is used to bias the Lin–Kernighan heuristic algorithm to pick edges are that more likely to be in the tour with lowest cost. Note that our approach is tightly connected to spectral methods for graphs [21]. Although the overall computational cost of computing the tours using the Procrustes solution is $O(n^3)$ – compared to $O(n^{2.2})$ in the case of traditional Lin–Kernighan heuristic [15] –, the overall Lin–Kernighan heuristic computation converges faster (in fewer iterations) than the 1-tree based approach.

Our paper is organized as follows: We start with the mathematical formulation of the TSP in Section 2 and construction of dynamical systems on the manifold of orthogonal matrices that converge to Hamiltonian cycles in Section 3. These isospectral flows are found to converge to undesirable local minima. In Section 4, we then describe the standard Lin–Kernighan heuristic as well as techniques to limit the search space using α-nearness, which is based on minimum spanning trees. We follow this discussion with a Procrustes-based approach for biasing the Lin–Kernighan heuristic called P-nearness (that is inspired by isospectral flows) in Section 5. Numerical results are presented in Section 6. We conclude with future work in Section 7.

2 The traveling salesman problem

Given a list of n cities $\{C_1, C_2, \ldots, C_n\}$ and the associated distances between cities C_i and C_j, denoted by d_{ij}, the TSP aims to find an ordering σ of $\{1, 2, \ldots, n\}$ such that the tour

cost, given by

$$c = \sum_{i=1}^{n-1} d_{\sigma(i),\sigma(i+1)} + d_{\sigma(n),\sigma(1)}, \tag{1}$$

is minimized. For the Euclidean TSP, for instance, $d_{ij} = \|x_i - x_j\|_2$, where $x_i \in \mathbb{R}^d$ is the position of C_i. In general, however, the distance matrix $D = (d_{ij})$ does not have to be symmetric. The ordering σ can be represented as a unique permutation matrix P. Note, however, that due to the cyclic symmetry multiple orderings – corresponding to different permutation matrices – have the same cost.

There are several equivalent ways to define the cost function of the TSP. We restrict ourselves to the trace[1] formulation proposed in [22]. Let $\mathcal{P}_n$ denote the set of all $n \times n$ permutation matrices, then the TSP can be written as a combinatorial optimization problem of the form

$$\min_{P \in \mathcal{P}_n} \operatorname{tr}\left(A^T P^T B P\right), \tag{2}$$

where $A = D$ and $B = T$. Here, T is defined to be the adjacency matrix of the cycle graph of length n. For the symmetric TSP, either the adjacency matrix of the directed cycle graph or the undirected cycle graph can be used, i.e.,

$$T_{\text{dir}} = \begin{pmatrix} 0 & 1 & & & \\ & 0 & 1 & & \\ & & \ddots & \ddots & \\ & & & 0 & 1 \\ 1 & & & & 0 \end{pmatrix} \quad \text{or} \quad T_{\text{undir}} = \begin{pmatrix} 0 & 1 & & & 1 \\ 1 & 0 & 1 & & \\ & \ddots & \ddots & \ddots & \\ & & 1 & 0 & 1 \\ 1 & & & 1 & 0 \end{pmatrix}.$$

The equivalence of (1) and (2) can be derived easily using the observation that $\tilde{T} := P^T T P$ is a permuted tour matrix, i.e., the (i, j) entry is 1 if the tour goes from city C_i to city C_j. Thus, for any permutation, $\operatorname{tr}(D^T \tilde{T}) = \sum_{i,j=1}^{n} d_{ij} \tilde{t}_{ij}$ is simply the sum of the distances associated with each edge. In what follows, we restrict our work to symmetric matrices; thus, we simply consider tour matrix $T = T_{\text{undir}}$ for undirected graphs.

The TSP can also be regarded as a special case of the general quadratic assignment problem (QAP) [23, 24], given by

$$\min_{P \in \mathcal{P}_n} \operatorname{tr}\left(A^T P^T B P + P^T C\right), \tag{3}$$

or a special case of the graph matching problem. The relationship between various combinatorial optimization problems is explored in Figure 1. In order to convert the minimization problem into the corresponding maximization problem, note that

$$\begin{aligned} \|A - P^T B P\|_F^2 &= \operatorname{tr}\left(A^T A\right) - 2\operatorname{tr}\left(A^T P^T B P\right) + \operatorname{tr}\left(B^T B\right) \\ &= \|A\|_F^2 - 2\operatorname{tr}\left(A^T P^T B P\right) + \|B\|_F^2. \end{aligned} \tag{4}$$

Thus, the norm is minimized if the trace is maximized and vice versa.

[1] The trace of a matrix $A \in \mathbb{R}^{n \times n}$ is defined to be the sum of all diagonal entries, i.e., $\operatorname{tr}(A) = \sum_{i=1}^{n} a_{ii}$.

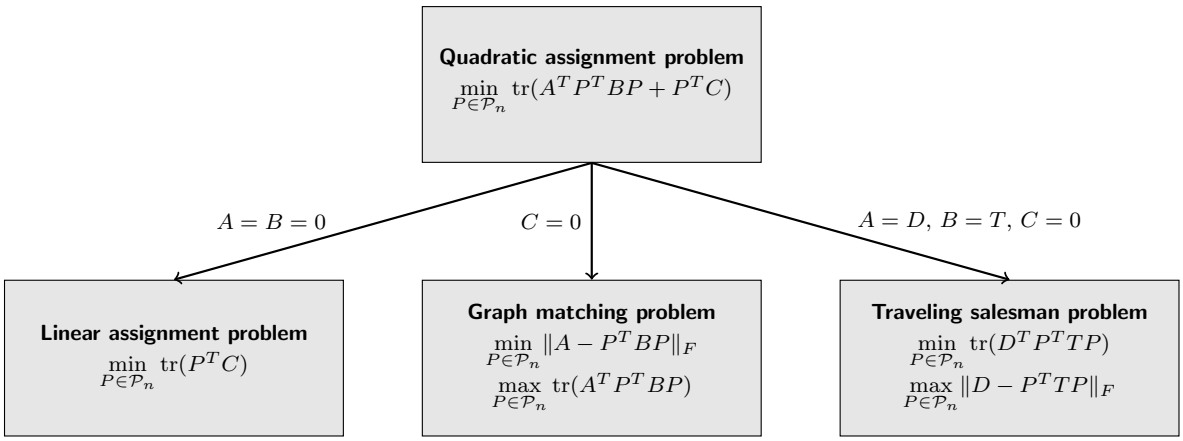
Quadratic assignment problem
$\min_{P \in \mathcal{P}_n} \operatorname{tr}(A^T P^T B P + P^T C)$
$A = B = 0$
$C = 0$
$A = D,\ B = T,\ C = 0$
Linear assignment problem
$\min_{P \in \mathcal{P}_n} \operatorname{tr}(P^T C)$
Graph matching problem
$\min_{P \in \mathcal{P}_n} \|A - P^T B P\|_F$
$\max_{P \in \mathcal{P}_n} \operatorname{tr}(A^T P^T B P)$
Traveling salesman problem
$\min_{P \in \mathcal{P}_n} \operatorname{tr}(D^T P^T T P)$
$\max_{P \in \mathcal{P}_n} \|D - P^T T P\|_F$

Figure 1: Relationships between the various combinatorial optimization problems.

Over the last decades, a plethora of heuristics has been developed to solve the TSP efficiently. In order to find a good approximation of the optimal tour, typically different global and local heuristics are combined. A very efficient and powerful approach is to construct an initial solution with the aid of greedy algorithms, for instance the nearest neighbor heuristic, and to improve the solution successively using local heuristics such as k-opt. One of the best available TSP solvers is Helsgaun's LKH software [17]. More details will be provided in Section 4.

3 Dynamical systems approach

In this section, we construct a dynamical systems approach for computing optimal tours for the TSP. In particular, we use matrix differential equations defined on the manifold of orthogonal matrices. As mentioned in the previous section, solutions of the TSP can be represented as permutation matrices. Permutation matrices lie on the manifold of orthogonal matrices. Our goal is to construct flows that minimize the TSP cost as they evolve. Note that gradient flow methods were first used by Brockett to compute eigenvalues and to solve linear programming or least squares matching problems [25, 26] and later also for combinatorial optimization problems [22, 27, 28]. We will formulate multiple cost functions for constructing gradient flows for the TSP.

3.1 Gradient flows for orthogonal matrices

Let $\mathcal{O}_n = \left\{ P \in \mathbb{R}^{n \times n} \mid P^T P = I \right\} \supset \mathcal{P}_n$ be the set of all $n \times n$ orthogonal matrices. We now consider the orthogonal relaxation of the combinatorial optimization problem (2), given by

$$\min_{P \in \mathcal{O}_n} \operatorname{tr}\left(A^T P^T B P\right). \tag{5}$$

This minimization problem can be solved using a steepest descent method on the manifold of orthogonal matrices. Given a cost function F, the gradient flow is defined as

$$\dot{P} = -\nabla F(P), \tag{6}$$

which is a matrix differential equation evolving on the manifold of orthogonal matrices. That is, starting with an orthogonal matrix P, the trajectory remains for all time in $\mathcal{O}_n$. Let $[A, B] = AB - BA$ be the standard Lie bracket and $\{A, B\} = A^T B - B^T A$ the generalized Lie bracket. The gradient of a function F defined on the manifold of orthogonal matrices is

$$\nabla F(P) = F_P - P F_P^T P = P\{P, F_P\}, \tag{7}$$

where F_P is the matrix of partial derivatives [29], i.e., $(F_P)_{ij} = \frac{\partial F}{\partial P_{ij}}$.

Lemma 3.1. *For the cost function $F(P) = \operatorname{tr}\left(A^T P^T B P\right)$, we obtain*

$$\nabla F(P) = P\left(\left\{P^T B P, A\right\} + \left\{P^T B^T P, A^T\right\}\right).$$

Proof. Since $\frac{\partial F}{\partial P} = BPA^T + B^T PA$, see [30], (7) leads to $\nabla F(P) = P\{P, BPA^T + B^T PA\}$, which can be rewritten as above. □

This is a generalization of the matrix flow defined in [27] for the symmetric graph matching problem. If A and B are symmetric, this can be simplified to

$$\nabla F(P) = 2P\left[P^T B P, A\right]. \tag{8}$$

Since the optimal solution of this optimization problem is in general not a permutation matrix, Zavlanos and Pappas [27] use a second term which penalizes nonnegative entries. Note that the set of permutation matrices is the intersection of the sets of orthogonal and nonnegative matrices. In order to force the gradient flow to converge to a permutation matrix, a cubic penalty function is used.

Lemma 3.2. *Let $\circ$ denote the Hadamard or element-wise product of two matrices. For the penalty function $G(P) = \frac{1}{3}\operatorname{tr}\left(P^T\left(P - (P \circ P)\right)\right) = \frac{1}{3}n - \frac{1}{3}\sum_{i,j=1}^{n} p_{ij}^3$, the gradient is given by*

$$\nabla G(P) = P\left((P \circ P)^T P - P^T (P \circ P)\right). \tag{9}$$

Proof. Note that $G_P = -(P \circ P)$. Using (7) results in the above gradient. □

By combining the two functions F and G, it is possible to compute a permutation matrix which is close to the optimal orthogonal solution. In [27], the steady state solution of the superimposed gradient flows for F and G, given by

$$\begin{aligned}\dot{P} = &-(1-k)P\left(\left\{P^T B P, A\right\} + \left\{P^T B^T P, A^T\right\}\right)\\ &- k\,P\left((P \circ P)^T P - P^T (P \circ P)\right),\end{aligned} \tag{10}$$

is computed for $k = 0$, then the parameter k is set to a value sufficiently close to 1 so that the flow converges to a permutation matrix. Another approach is to apply a homotopy-based method where k is the continuation parameter which is gradually increased until the solution is close to a permutation matrix.

Example 3.3. In order to illustrate the gradient flow approach, let us consider a simple TSP with 10 cities. Using (10), we obtain the results shown in Figure 2. For this system, the dynamical system converges to the optimal tour.

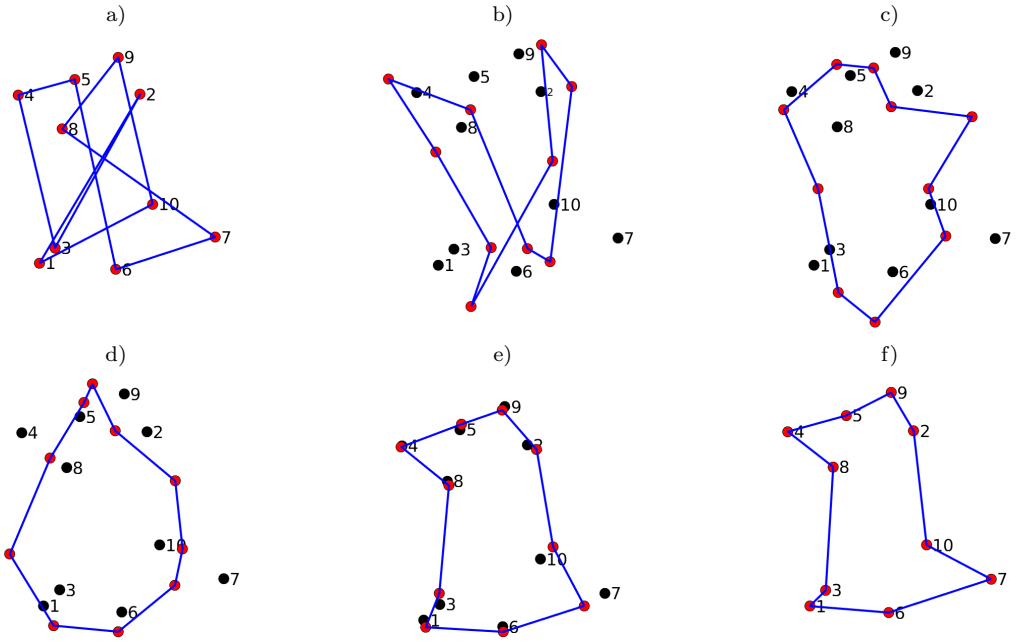

Figure 2: Traveling salesman problem with 10 cities solved using the gradient flow (10). The original positions of the cities are shown in black, the positions transformed by the orthogonal matrix P in red. a) Initial trivial tour given by $\sigma = (1, \ldots, 10)$. b–c) Intermediate solutions. d) Optimal orthogonal matrix for $k = 0$. e) Convergence to a permutation matrix for increasing k. f) Final solution. The initial tour was transformed into the optimal tour by the gradient flow.

It is also possible to write the relaxed TSP as a constrained optimization problem of the form

$$\begin{aligned} &\min\ \operatorname{tr}\left(A^T P^T B P\right), \\ &s.t.\ \ \tfrac{1}{3}\operatorname{tr}\left(P^T\left(P - (P \circ P)\right)\right) = 0 \end{aligned}$$

and solve the resulting system of differential equations

$$\begin{aligned} \dot{P} &= -P\left(\left\{P^T B P, A\right\} + \left\{P^T B^T P, A^T\right\}\right) - \lambda P\left((P \circ P)^T P - P^T (P \circ P)\right), \\ \dot{\lambda} &= \tfrac{1}{3}\operatorname{tr}\left(P^T\left(P - (P \circ P)\right)\right). \end{aligned}$$

That is, we perform a gradient descent for the cost function and a gradient ascent for the Lagrange multiplier, as described in [31] for general constrained optimization problems.

Remark 3.4. In related work, Wen and Yin [28] recently proposed a gradient descent method for optimization problems with orthogonality constraints. Their method can be interpreted as a structure-preserving integration scheme for the corresponding matrix differential equation. Let $M_k = F_{P_k} P_k^T - P_k F_{P_k}^T$, then the trapezoidal rule for (6) can be written as

$$P_{k+1} = \left(I + \tfrac{h}{2} M_{k+1}\right)^{-1} \left(I - \tfrac{h}{2} M_k\right) P_k,$$

while the update equation in [28] is defined as

$$P_{k+1} = \underbrace{\left(I + \tfrac{h}{2} M_k\right)^{-1} \left(I - \tfrac{h}{2} M_k\right)}_{Q_k} P_k. \tag{11}$$

Since M_k is skew-symmetric, the matrix Q_k is orthogonal (Cayley transform) and the new solution P_{k+1} automatically satisfies the orthogonality constraints. Conventional integration schemes will in general not preserve these constraints. The approach is demonstrated on an instance of the QAP.

For the TSP, there are several different ways to define a cost function. Similarly, various penalty functions could be used to force the gradient flow to converge to a permutation matrix. It turns out that the cubic penalty function from Lemma 3.2 often converges to unfavorable solutions, in particular for large n. Since the number of solutions grows exponentially, several different initial conditions have to be used to find a good solution of the TSP.

3.2 Gradient flows for tour matrices

Instead of forcing the gradient flow to converge to a permutation matrix, an alternative approach is to define a cost function in such a way that the flow converges to a permutation of the initial tour matrix T. For the symmetric case (8), Brockett [26] introduces a change of variables, given by $H = P^T B P$. The resulting double bracket flow, $\dot{H} = 2\,[H, [H, A]]$, now evolves in the space of symmetric matrices and is only quadratic in H.

Assuming that the objective function $F(P)$ can be rewritten as a function $F(H)$, we now want to derive a gradient flow for the new variable H. Here, we do not assume that B is symmetric. For the TSP, again $A = D$ and $B = T$. With the aforementioned transformation, the cost function for the relaxed TSP from Lemma 3.1 can be written as

$$\min_{H \in \mathcal{T}_n} \operatorname{tr}\left(A^T H\right), \tag{12}$$

where $\mathcal{T}_n = \{P^T T P \mid P \in \mathcal{O}_n\}$.

Remark 3.5. For directed cycle graphs, i.e., $B = T_{\text{dir}}$, the non-relaxed version of (12) is identical to the cost of the LAP, with the difference that here the set of feasible solutions is constrained to a subset of $\mathcal{P}_n$, which makes the problem NP-hard.

Theorem 3.6. *Let $F(H)$ be a given cost function, then the gradient flow is given by*

$$\dot{H} = -\left[H, \{H, F_H\} + \left\{H^T, F_H^T\right\}\right].$$

Proof. Since $H = P^T BP$, using (6) and (7) we obtain

$$\dot{H} = \dot{P}^T BP + P^T B\dot{P} = -\left[H, \{P, F_P\}\right].$$

Applying the chain rule, this leads to

$$\begin{aligned}\frac{\partial F}{\partial P_{ij}} &= \operatorname{tr}\left(\left(\frac{\partial F}{\partial H}\right)^T \frac{\partial H}{\partial P_{ij}}\right) \\ &= \operatorname{tr}\left(F_H^T \left(P^T B J^{ij} + J^{ji} BP\right)\right) \\ &= \left(B^T P F_H + BP F_H^T\right)_{ij},\end{aligned}$$

where $J^{ij} \in \mathbb{R}^{n\times n}$ is a single-entry matrix [30], i.e., $(J^{ij})_{kl} = \delta_{ik}\delta_{jl}$. It follows that

$$F_P = B^T P F_H + BP F_H^T. \tag{13}$$

Inserting this into the equation for $\dot{H}$ concludes the proof. □

For the cost function $F(H) = \operatorname{tr}(A^T H)$, we simply obtain $F_H = A$. With the aid of Theorem 3.6, we can then compute the corresponding gradient flow. In addition to the cost function, a penalty function has to be used to find an admissible solution. One possibility would be to penalize negative entries, as described above. Furthermore, linear equality constraints that force the flow to converge to a matrix H with row and column sums equal to two could be added. We use the penalty function

$$G(H) = \|H - H \circ H\|_F,$$

which penalizes entries that are not zero or one. We then obtain

$$G_H = 2(H - H \circ H) \circ (E - 2H). \tag{14}$$

Here, E denotes the matrix of ones. Thus, the overall flow is given by

$$\begin{aligned}\dot{H} = &-(1-k)\left[H, \{H, F_H\} + \{H^T, F_H^T\}\right] \\ &- k\ \left[H, \{H, G_H\} + \{H^T, G_H^T\}\right],\end{aligned} \tag{15}$$

where F_H and G_H are as derived above.

Example 3.7. Let us illustrate the gradient flow with the same TSP with 10 cities as in Example 3.3. Using (15), we obtain the results shown in Figure 3. We plot only the entries of the H matrix that are greater than zero. The dynamical system converges to a tour that is slightly longer than the optimal tour.

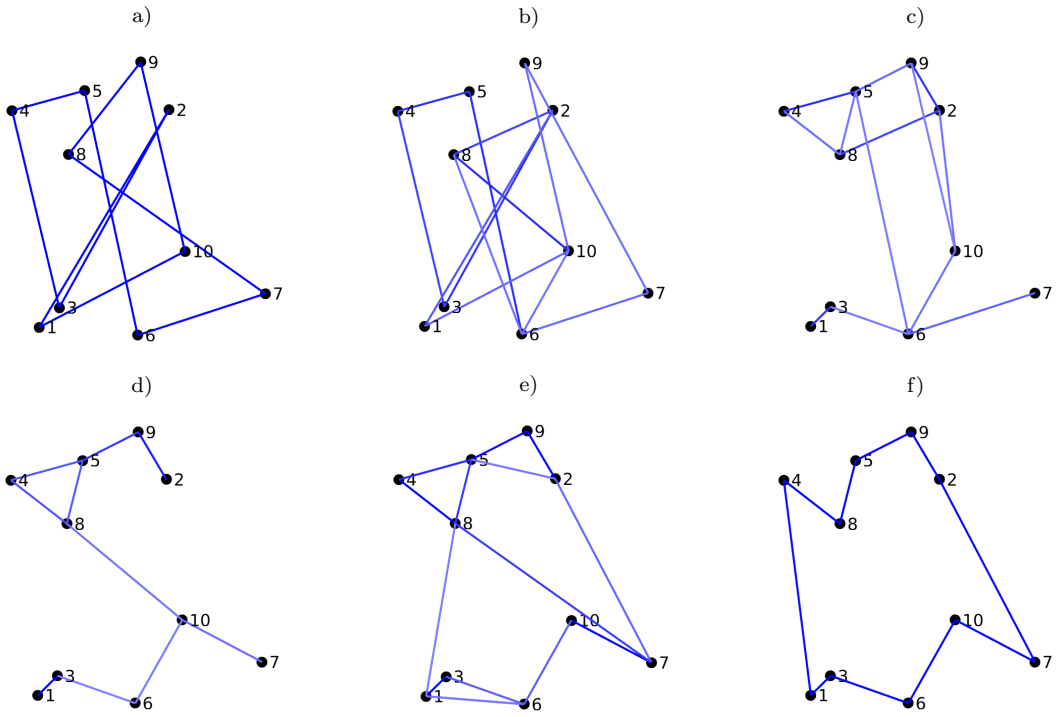

Figure 3: Traveling salesman problem with 10 cities solved using the gradient flow (15). a) Initial trivial tour T. b–c) Intermediate solutions. d) Optimal H matrix for $k = 0$. e) Convergence to the tour matrix for increasing k. f) Final solution. The initial tour matrix was transformed into a new shorter tour.

3.3 A structure-preserving integrator

Given an arbitrary cost function $F(H)$, we now derive a structure-preserving integrator for the flow evolving in the space of symmetric matrices using (11) and (13). Note that

$$\begin{aligned} P_k^T M_k P_k &= P_k^T \left(F_{P_k} P_k^T - P_k F_{P_k}^T \right) P_k \\ &= P_k^T \left(B^T P_k F_{H_k} + B P_k F_{H_k}^T \right) - \left(F_{H_k}^T P_k^T B + F_{H_k} P_k^T B^T \right) P_k \\ &= H_k^T F_{H_k} + H_k F_{H_k}^T - F_{H_k}^T H_k - F_{H_k} H_k^T \\ &= \left\{ H_k, F_{H_k} \right\} + \left\{ H_k^T, F_{H_k}^T \right\}. \end{aligned}$$

Let $N_k = \left\{ H_k, F_{H_k} \right\} + \left\{ H_k^T, F_{H_k}^T \right\}$, then

$$\begin{aligned} P_{k+1} &= P_k \left(P_k^T \left(I + \tfrac{h}{2} M_k \right) P_k \right)^{-1} \left(P_k^T \left(I - \tfrac{h}{2} M_k \right) P_k \right) \\ &= P_k \left(I + \tfrac{h}{2} N_k \right)^{-1} \left(I - \tfrac{h}{2} N_k \right) \end{aligned}$$

and

$$H_{k+1} = P_{k+1}^T B P_{k+1} = \underbrace{\left(I - \tfrac{h}{2}N_k\right)^{-1}\left(I + \tfrac{h}{2}N_k\right)}_{R_k^T} H_k \underbrace{\left(I + \tfrac{h}{2}N_k\right)^{-1}\left(I - \tfrac{h}{2}N_k\right)}_{R_k}.$$

Since the matrix R_k is orthogonal, H_{k+1} and H_k are similar and the eigenvalues are preserved. Using this structure-preserving integrator, the matrix differential equation can be solved efficiently. An implicit integrator for the matrix differential equation would require the computation of the $n^2 \times n^2$ Jacobian. This would be prohibitively inefficient in terms of both CPU time and memory usage for larger problems. Before moving to the Lin–Kernighan heuristic, we would like to point out that the above flow formulation is also applicable to the QAP described in Section 2. In other words, one can construct flows for the QAP that are frequently comparable to competing state-of-the-art approaches [32].

4 The Lin–Kernighan heuristic

The Lin–Kernighan heuristic is a popular heuristic approach for the TSP introduced in [15]. Starting from a tour, the approach progresses by extracting edges from the tour and replacing them with new edges, while maintaining the Hamiltonian cycle constraint. If k edges in the tour are replaced, this is known as the k-opt move [33]. To reduce the search space, the algorithm uses minimum spanning trees [19, 20] to identify edges that are more likely to be in the tour. The algorithm has found great success [17, 18] on large instances, see [1] for more details. Note that this algorithm has been extended to generalized TSPs [34] and clustered TSPs [35].

The LKH package [17, 18] offers different heuristics to compute an initial tour. The standard method is to choose one node at random and to iteratively add edges based on computed α-nearness values and related candidate sets until a tour is found. As described below, LKH uses minimum spanning trees to compute α-nearness. When an initial tour has been found, LKH improves it using local heuristics. A very popular and efficient heuristic is k-opt. The simplest version, 2-opt, removes two edges of the tour and reconnects the subtours as shown in Figure 4. If the resulting tour is shorter than the original tour, the step is accepted and rejected otherwise. Accordingly, 3-opt removes three edges of the current tour, reconnects the subtours and picks the shortest tour.

4.1 Candidate sets and α-nearness

LKH uses k-opt with varying k. The basic move is a sequential 5-opt step. In order to limit the search space and to increase the efficiency of k-opt moves, candidate sets that contain promising edges are computed for all cities. Methods to construct candidate sets for large TSPs have to be efficient in terms of both CPU time and memory usage. The standard approach implemented in LKH, which is called α-nearness, is based on minimum spanning trees or, to be more precise, on 1-trees, a slight variant of minimum spanning trees.

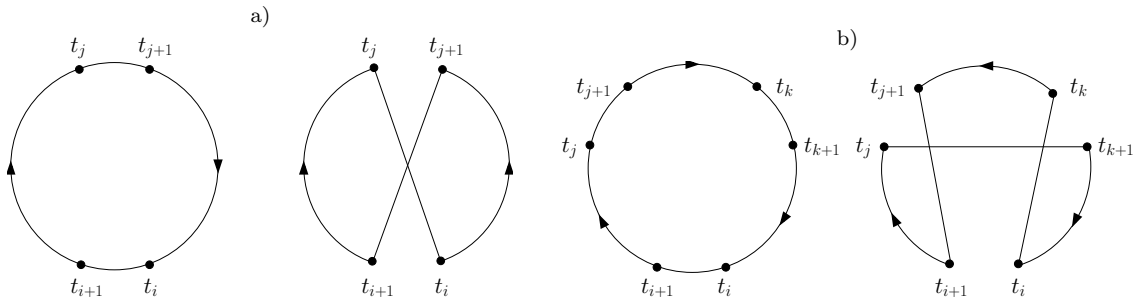

Figure 4: 2-opt and 3-opt move.

Definition 4.1. *Let $G = (V, E)$ be a graph with vertices V and edges E. A 1-tree for G is defined to be a spanning tree for the vertices $V \setminus \{v_1\}$ plus two additional edges $e \in E$ incident to vertex v_1.*

A 1-tree with minimum weight is called a minimum 1-tree. Note that every tour is a 1-tree with the additional property that the degree of each vertex is two. The minimum 1-tree usually contains already many edges of the shortest tour. The definition of α-nearness is based on a sensitivity analysis using 1-trees. The α-nearness value for the cities C_i and C_j is, roughly speaking, the difference between a minimum 1-tree and a 1-tree that is required to contain edge (v_i, v_j). That is, if an edge belongs to the minimum 1-tree, then the α-nearness value is 0 and the edge is assigned a high probability of being part of the shortest tour.

For each city, the candidate set is then defined to be the set of the m incident edges with the lowest α-nearness values. The candidate sets are used to limit and also to direct the search. Candidate sets based only on the distance between cities are typically not connected (see also Figure 7) and convergence to a good solution of the TSP might be slow. It was shown by Stewart [36] that minimum spanning trees, which are by definition always connected, can be used to increase the efficiency of local heuristics.

4.2 Subgradient optimization

In order to improve the α-nearness values, a subgradient optimization method which changes the original distance matrix D in such a way that the degree of almost all vertices of the optimized 1-tree is 2 is applied. The entries of the new distance matrix, $\widetilde{D}(\pi)$, are computed as

$$\widetilde{d}_{ij}(\pi) = d_{ij} + \pi_i + \pi_j.$$

The π values, sometimes called penalties, change the distances between the cities. The basic idea is to make edges incident to vertices with a low degree shorter and edges incident to vertices with a high degree longer so that the resulting 1-tree is close to a tour. This transformation of the distance matrix does not change the shortest tour and leads to significantly improved α-nearness values [17]. This method can also be used to compute a lower bound which is in general very close to the optimal tour length [19, 20]. Figure 5 shows the impact of the subgradient optimization. In the example, most of the edges of the optimal tour are

already present in the optimized 1-tree. For a more detailed description of α-nearness, 1-trees, and the subgradient optimization scheme, we refer to [19, 20, 17].

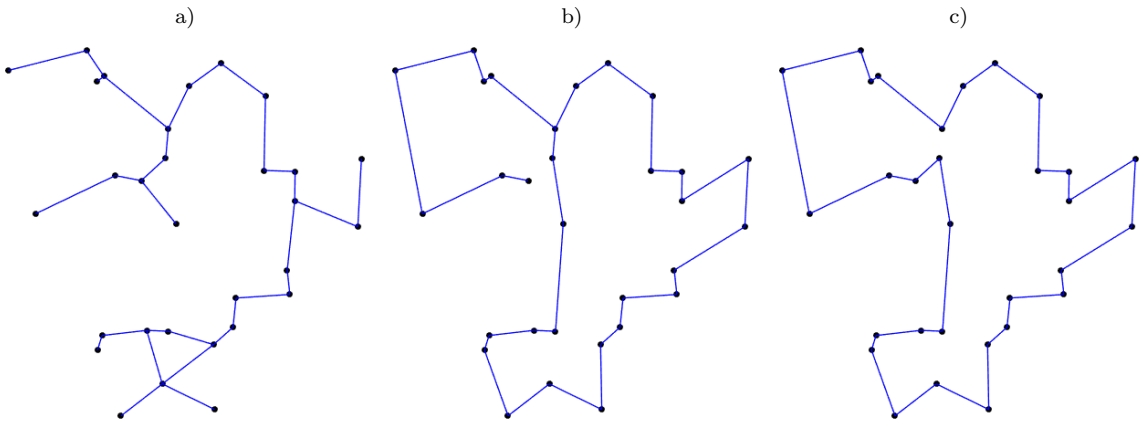

Figure 5: Impact of the subgradient optimization. a) Minimal 1-tree of the original distance matrix. b) Minimal 1-tree of the transformed distance matrix. c) Shortest tour.

5 Procrustes-based Lin–Kernighan heuristic

In this section, we will propose a method to compute candidate sets based on the relaxed problem (5) or (12), respectively. The solution, which is given by the solution of a related problem, can be computed analytically. Note that the solutions computed using this approach are optimal solutions of flows (for P and H) described in previous sections.

5.1 The two-sided orthogonal Procrustes problem

Let A and B be two symmetric $n \times n$ matrices. Then

$$\min_{P \in \mathcal{O}_n} \|A - P^T B P\|_F \tag{16}$$

is called the *two-sided orthogonal Procrustes problem*. As shown in (4), cost function (5) is minimized if the cost function $\|A - P^T B P\|_F$ of the Procrustes problem is maximized and vice versa. Since $\mathcal{P}_n \subset \mathcal{O}_n$, the cost of the orthogonal matrix is always lower than (or equal to if the matrices A and B are permutation-similar) the cost of the permutation matrix.

Theorem 5.1. *Given two symmetric matrices A and B, whose eigenvalues are distinct, let $A = V_A \Lambda_A V_A^T$ and $B = V_B \Lambda_B V_B^T$ be eigendecompositions, with $\Lambda_A = \operatorname{diag}\left(\lambda_A^{(1)}, \ldots, \lambda_A^{(n)}\right)$, $\Lambda_B = \operatorname{diag}\left(\lambda_B^{(1)}, \ldots, \lambda_B^{(n)}\right)$, and $\lambda_A^{(1)} \geq \cdots \geq \lambda_A^{(n)}$ as well as $\lambda_B^{(1)} \geq \cdots \geq \lambda_B^{(n)}$. Then every orthogonal matrix P^* which minimizes (16) has the form*

$$P^* = V_B S V_A^T,$$

where $S = \operatorname{diag}(\pm 1, \ldots, \pm 1)$.

A proof of this theorem can be found in [37], for example. If the eigenvalues of A and B are distinct, then there exist 2^n different solutions with the same cost. If one or both of the matrices possess repeated eigenvalues, then the eigenvectors in the matrices V_A and V_B are determined only up to basis rotations, which further increases the solution space.

The theorem states that in order to minimize the cost function in (16), the eigenvalues and corresponding eigenvectors have to be sorted both in either increasing or decreasing order. On the other hand, from the proof of the theorem it can be seen that in order to compute the solution of (5), the eigenvalues and eigenvectors of A and B, respectively, have to be sorted in opposite order.

Thus the optimal solution which minimizes the relaxed TSP cost function (5) can be obtained using the matrix $P^* = V_T V_D^T$ [38, 39], where the eigenvectors in the two matrices are sorted with respect to increasing eigenvalues of T and decreasing eigenvalues of D or vice versa. Define $T^* = P^{*T} T P^* = V_D \Lambda_T V_D^T$ to be the solution of the two-sided orthogonal Procrustes problem. Note that T^* is also the optimal solution of the gradient flow (15) for $k = 0$. Roughly speaking, T^* can be interpreted as a *continuous* solution of the relaxed TSP where the entry t^*_{ij} describes the strength of edge (i, j).

Example 5.2. Let us consider again the TSP from Example 3.3 and Example 3.7. Figure 6 shows the difference between the optimal solution of the TSP and the optimal solution of the Procrustes problem. We use a linear interpolation between blue (large t^*_{ij} value) and white (small t^*_{ij} value). Note that this is the same matrix as in Figure 3d, with the difference that we are plotting a few more edges here to illustrate P-nearness. Clearly, some edges of the optimal tour are already visible in Figure 6b, for example $(2, 9)$ and $(5, 9)$, while other edges such as $(3, 10)$ or $(2, 5)$ have a much lower weight. For city 10, different choices exist, $(2, 10)$, $(6, 10)$, $(7, 10)$, and $(8, 10)$, for instance, have a high probability of being part of the shortest tour.

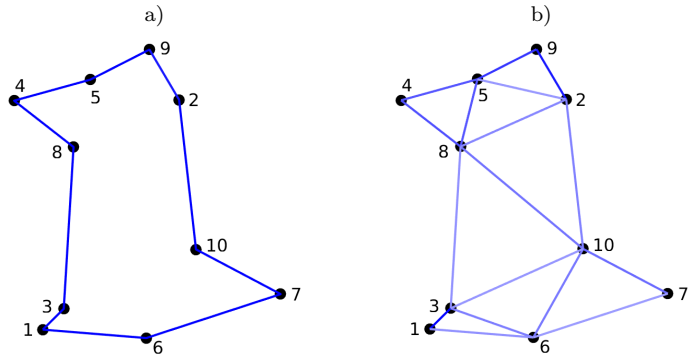

Figure 6: Solution of a TSP with 10 cities. a) Optimal permutation matrix. b) Optimal orthogonal matrix.

We want to use the solution of the Procrustes problem to limit the search space and to improve the efficiency of local heuristics. The aim is to bias k-opt in such a way that edges with high edge strengths are included with a high probability. We compute candidate sets

based on the entries t^*_{ij} of the matrix T^* and call this approach P-nearness. That is, for each city, we pick the cities with the largest entries t^*_{ij}.

In Figure 7, we show the edges computed using the Procrustes solution. In particular, for two random TSP instances (50 city and 100 city examples), we show the shortest edges in the left-most column and the edges from the Procrustes solution in the right-most column. The optimal tour is plotted in the middle column. It is evident from the figure that the Procrustes solution T^* tends to capture most of the edges in the optimal tour. Note that, while the α-nearness values can be computed in $O(n^2)$ [19], the computation of the P-nearness values is $O(n^3)$ [17, 18].

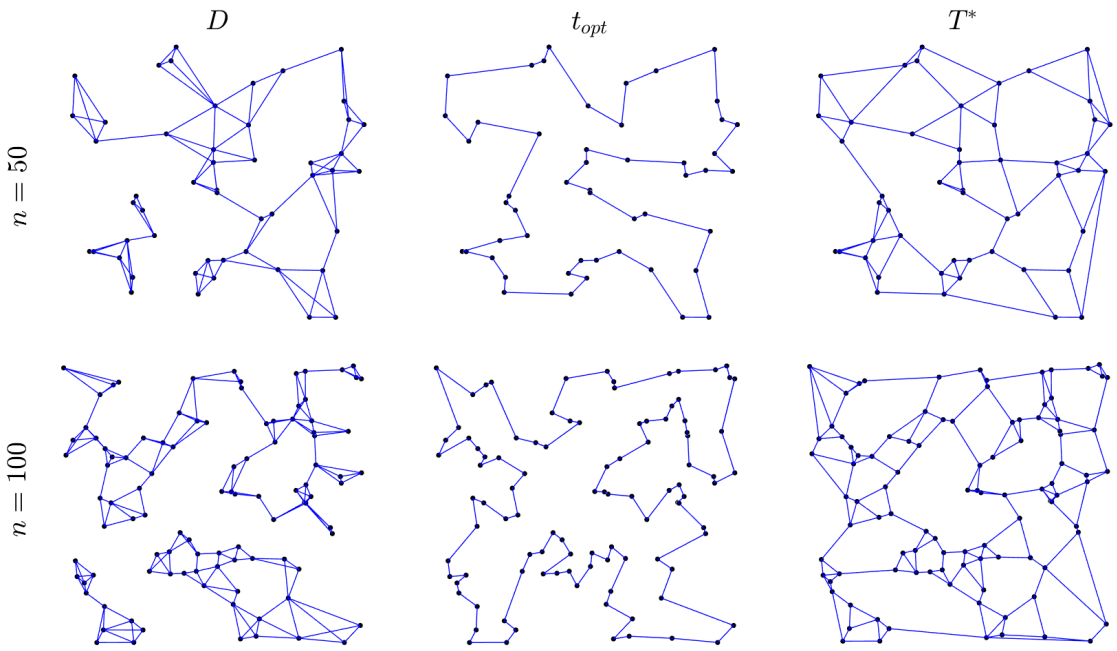

Figure 7: Illustration of P-nearness for random TSP instances of size 50 and 100. The left column contains the edges with shortest distance, the center column has the optimal tour for the instances, and the right column contains the edges with the highest P-nearness values for each city. For each city, we plotted the three edges with the highest nearness values.

5.2 Improving the Procrustes solution

We now describe our approach to obtain better P-nearness values than those computed from the solution of the Procrustes problem alone. The principal idea behind our method is to construct a homotopy between the original TSP distance matrix D and the solution of the Procrustes solution. Intuitively, one would desire the candidate sets to include 'several' short edges and a 'few' long edges. Note that simply picking the shortest edges from D gives rise to greedy solutions that are usually not competitive since they require the addition of long edges to complete the tours [1].

We find that the Procrustes solution T^* tends to select too many long edges (as shown in Figure 8a). If we use the entries of T^* to bias the Lin–Kernighan heuristic, the solutions are found to be less competitive than the standard approach. An effort to reduce the number of long edges is equivalent to making the computed solution "greedy" by picking edges based on the distance matrix D. Thus, we construct a homotopy $\tilde{H}$ of the form,

$$\tilde{H} = T^* - \lambda D.$$

The candidate sets for varying λ are shown for an example TSP instance in Figure 8. To find the optimal λ we use ideas from graph clustering, where one computes the existence of disconnected clusters in the graph. In particular, we increase λ until the graph of candidate sets is almost disconnected (separated into clusters). This optimal λ is found by either marching in λ or using a bisection approach.

There are multiple ways that one can compute the connectedness of a graph. In particular, one can use a depth-first search based approach [40] or perform computations on the graph Laplacian [41, 42, 43]. The rank of the graph Laplacian matrix is related to the number of connected components in the graph [43]. In our work, we pick the graph Laplacian approach for computing connected components in the graph (by looking at the multiplicity of the zero eigenvalue). Note that these computations can also be performed in the distributed setting [44]. If varying λ does not give rise to a disconnected graph, we set $\lambda = 1$. Alternatively, one can use the D matrix (in place of $T^* - \lambda D$) to bias the k-opt moves in the Lin–Kernighan heuristic. If the candidate sets based on distance only are connected, this typically implies that k-opt converges quickly to the shortest tour.

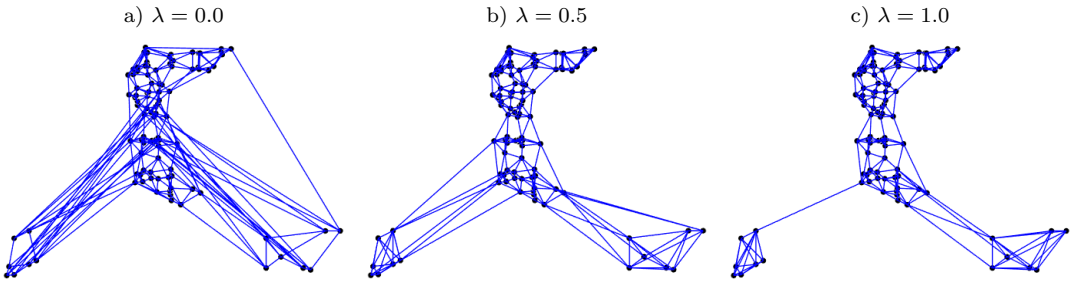

Figure 8: Homotopy of $T^* - \lambda D$.

6 Results

To compare different candidate sets or methods to bias k-opt, Helsgaun computes the average rank of the edges which form the shortest tour [17]. The optimal average rank is 1.5, all edges belonging to the shortest tour have either rank 1 or 2. We found that the average rank is in general not a good metric for the quality of the nearness values or candidate sets. Although the average rank of the Procrustes solution is typically much higher than the average rank of the α-nearness values, k-opt often converges faster to the shortest tour.

In order to compare α-nearness and P-nearness, we compute tours using Helsgaun's LKH package. For each LKH run, we generate the candidate sets based on the α-nearness and P-nearness values. Starting from initial tours computed using α-nearness and P-nearness, respectively, we compare the resulting tour lengths after a fixed number of k-opt steps.

In Table 1, we compare 22 well-known instances of the TSP from the TSPLIB database [45]. The size of the candidate sets in these computations is fixed, we compute 5 candidates for each city using α-nearness or P-nearness, respectively. Starting from a random initial tour that is generated from the respective candidate sets, we perform a fixed number of $8n$ k-opt moves, where n is the number of cities. We find that in this setting, the P-nearness based approach converges faster than α-nearness. For example, after $8n$ steps, P-nearness based LKH converges to lower cost values in 18 of the instances when compared to α-nearness based LKH. Moreover, we ran 50 random TSP instances of size 1000 (cities) and found that P-nearness had lower tour costs after a fixed number of k-opt moves in 31 of the instances, hence resulting is better solutions in 62% of the instances. Note that if we run both, α-nearness and P-nearness based LKH, to convergence, both methods compute the best known optimal tours in these instances.

7 Conclusion and future work

In this work, we explored the use of continuous relaxations and dynamical systems theory for constructing algorithms for the TSP. Our approach aimed to exploit the observation that the solution of the TSP can be represented as a permutation matrix which lies on the manifold of orthogonal matrices. In the first part of this manuscript, we constructed a dynamical system on the manifold of orthogonal matrices that converges to solutions of the TSP. We also explored the construction of gradient flows for tour matrices in Section 3.2. We found that although the dynamical systems approach is elegant, it often converges to local optima.

Inspired by the dynamical systems approach, we then constructed a Procrustes based approach that computes an orthogonal matrix that minimizes the TSP cost. Our approach was based on the computation of the solution of the two-sided orthogonal Procrustes problem. This solution is based on the eigendecomposition of the tour and distance matrices. We then constructed a homotopy with the distance matrix that is then used to bias the popular Lin–Kernighan heuristic.

The candidate sets constructed from the homotopy are found to give faster convergence than minimum spanning tree (1-tree) based approaches. Our algorithm was implemented in the LKH software framework and demonstrated on multiple TSPLIB and random TSP instances.

Future work includes the generalization of our proposed approach to QAPs. The aim is to develop efficient heuristics utilizing the results of the Procrustes problem and dynamical systems theory for solving strongly NP-hard problems. We are also exploring the use of subgradient optimization for improving the Procrustes solution which is expected to provide faster convergence rates for the TSP and related optimization problems.

Table 1: Comparison of α-nearness and P-nearness based Lin–Kernighan heuristic on TSPLIB instances. The size of the candidate sets is set to 5 per city and we stop the computations after $8n$ k-opt moves in LKH, where n is the number of cities. Out of the 22 TSPLIB instances, P-nearness computes a better solution in 18 of the cases.

TSP	α-nearness	P-nearness	improvement
d198	16540	16465	0.45 %
pcb442	50785	50832	-0.09 %
d493	36028	35023	2.79 %
u574	36984	36926	0.16 %
rat575	6796	6790	0.09 %
p654	35716	37039	-3.70 %
d657	49504	49158	0.70 %
u724	42295	41904	0.92 %
rat783	9054	8810	2.69 %
pr1002	261797	259810	0.76 %
u1060	224510	224552	-0.20 %
vm1084	244411	242573	0.75 %
pcb1173	56934	56915	0.03 %
d1291	53357	51610	3.27 %
rl1323	279810	275904	1.40 %
nrw1379	141510	67035	52.63 %
fl1400	21319	22775	-6.83 %
u1432	153213	153054	0.10 %
fl1577	28217	24357	13.68 %
d1655	95532	64837	32.13 %
u1817	58351	58213	0.24 %
rl1889	345475	340271	1.51 %

8 Acknowledgements

The authors thank Prof. Keld Helsgaun for discussions related to the Lin–Kernighan heuristic and his software and also Dr. Mirko Hessel-von Molo and Steffen Ridderbusch for discussions related to the approach.

References

[1] W.J. Cook. *In Pursuit of the Traveling Salesman: Mathematics at the Limits of Computation.* Princeton University Press, 2011.

[2] R. Agarwala, D.L. Applegate, D. Maglott, G.D. Schuler, and A.A. Schäffer. A fast and scalable radiation hybrid map construction and integration strategy. *Genome Research*, 10(3):350–364, 2000.

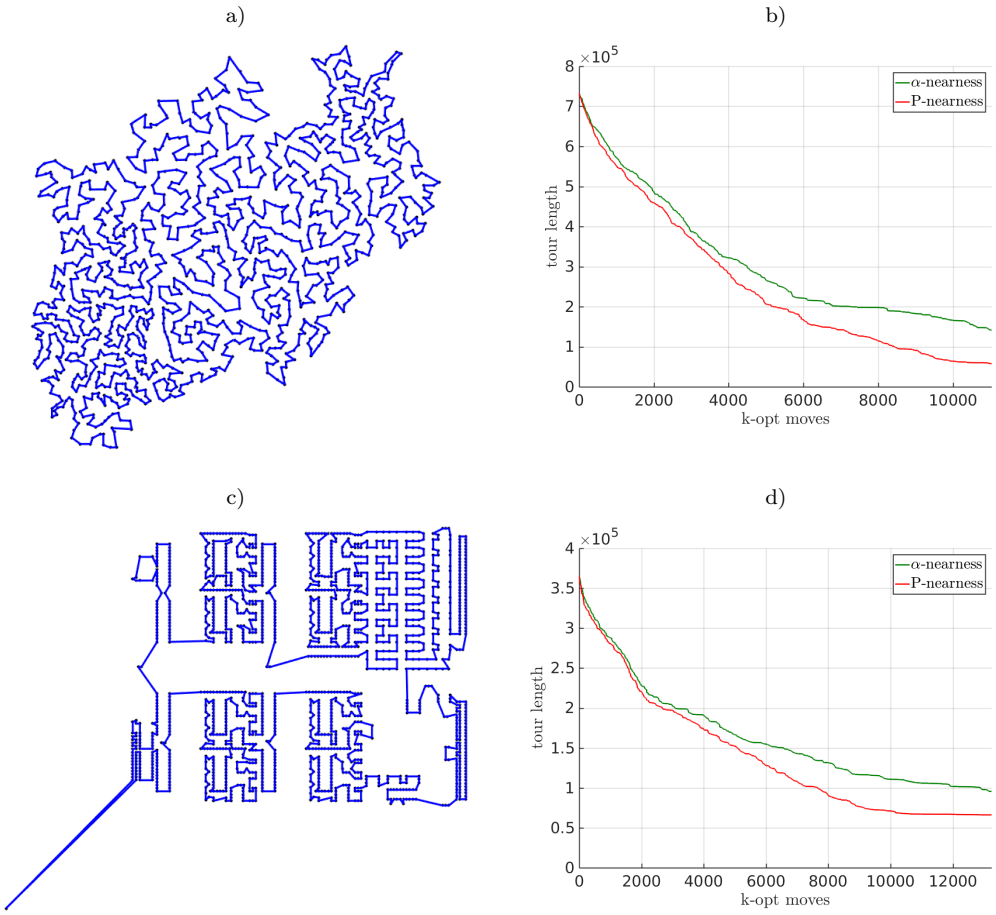

Figure 9: TSPLIB instances and tour lengths as a function of the number of k-opt moves. a–b) nrw1379. c–d) d1655.

[3] S. Carlson. Algorithm of the gods. *Scientific American*, 276:121–124, 1997.

[4] E. Kolemen and N.J. Kasdin. Optimal trajectory control of an occulter-based planet-finding telescope. *American Astronautical Society*, pages 07–037, 2007.

[5] M. Grötschel, M. Jünger, and G. Reinelt. Optimal control of plotting and drilling machines: a case study. *Mathematical Methods of Operations Research*, 35(1):61–84, 1991.

[6] B. Englot, T. Sahai, and I. Cohen. Efficient tracking and pursuit of moving targets by heuristic solution of the traveling salesman problem. In *52nd IEEE Conference on Decision and Control*, pages 3433–3438. IEEE, 2013.

[7] J.K. Lenstra and A.H.G. Kan. Complexity of vehicle routing and scheduling problems. *Networks*, 11(2):221–227, 1981.

[8] O. Goldschmidt, A. Laugier, and E.V. Olinick. Sonet/sdh ring assignment with capacity constraints. *Discrete Applied Mathematics*, 129(1):99–128, 2003.

[9] T. Sahai and S. Klus. Automatic learning of Bayesian networks, May 2015. US Patent App. 14/546,392.

[10] R.M. Karp. Reducibility among combinatorial problems. *50 Years of Integer Programming 1958–2008*, pages 219–241, 2010.

[11] M. Dorigo and L.M. Gambardella. Ant colony system: A cooperative learning approach to the traveling salesman problem. *Evolutionary Computation, IEEE Transactions on*, 1(1):53–66, 1997.

[12] M. Padberg and G. Rinaldi. A branch-and-cut algorithm for the resolution of large-scale symmetric traveling salesman problems. *SIAM review*, 33(1):60–100, 1991.

[13] D. Applegate, R. Bixby, W. Cook, and V. Chvátal. *On the solution of traveling salesman problems*. Rheinische Friedrich-Wilhelms-Universität Bonn, 1998.

[14] N. Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. Technical report, DTIC Document, 1976.

[15] S. Lin and B.W. Kernighan. An effective heuristic algorithm for the traveling-salesman problem. *Operations research*, 21(2):498–516, 1973.

[16] J.C. Gower and G.B. Dijksterhuis. *Procrustes problems*, volume 3. Oxford University Press New York, 2004.

[17] K. Helsgaun. An effective implementation of the Lin–Kernighan traveling salesman heuristic. Datalogiske Skrifter (Writings on Computer Science), No. 81, Roskilde University, 1998.

[18] Keld Helsgaun. *An effective implementation of K-opt moves for the Lin-Kernighan TSP heuristic*. PhD thesis, Roskilde University. Department of Computer Science, 2006.

[19] M. Held and R. M. Karp. The traveling-salesman problem and minimum spanning trees. *Operations Research*, 18(6):1138–1162, 1970.

[20] M. Held and R. M. Karp. The traveling-salesman problem and minimum spanning trees: Part II. *Mathematical Programming*, 1:6–25, 1971.

[21] S. Klus and T. Sahai. A spectral assignment approach for the graph isomorphism problem. *In Submission, http://arxiv.org/abs/1411.0969*, 2017.

[22] W. S. Wong. Matrix representation and gradient flows for NP-hard problems. *Journal of Optimization Theory and Applications*, 87:197–220, 1995.

[23] T. C. Koopmans and M. Beckmann. Assignment problems and the location of economic activities. *Econometrica*, 25(1):53–76, 1957.

[24] R. E. Burkard, E. Çela, P. M. Pardalos, and L. S. Pitsoulis. The quadratic assignment problem. In P. M. Pardalos and D.-Z. Du, editors, *Handbook of Numerical Optimization*. Kluwer Acedemic Publishers, 1998.

[25] R. W. Brockett. Least squares matching problems. *Linear Algebra and its Applications*, 122–124:761–777, 1989.

[26] R. W. Brockett. Dynamical systems that sort lists, diagonalize matrices and solve linear programming problems. *Linear Algebra and Its Applications*, 146:79–91, 1991.

[27] M. M. Zavlanos and G. Pappas. A dynamical systems approach to weighted graph matching. *Automatica*, 44:2817–2824, 2008.

[28] Z. Wen and W. Yin. A feasible method for optimization with orthogonality constraints. Technical report, Rice University, 2010.

[29] A. Edelman, T. A. Arias, and S. T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM Journal on Matrix Analysis and Applications*, 20(2):303–353, 1998.

[30] K. B. Petersen and M. S. Pedersen. The matrix cookbook, 2008.

[31] J. C. Platt and A. H. Barr. Constrained differential optimization for neural networks. Technical report, California Institute of Technology, 1988.

[32] M. Zaslavskiy. *Graph matching and its application in computer vision and bioinformatics*. PhD thesis, l'Ecole nationale superieure des mines de Paris, 2010.

[33] Keld Helsgaun. General k-opt submoves for the Lin–Kernighan TSP heuristic. *Mathematical Programming Computation*, 1(2-3):119–163, 2009.

[34] K. Helsgaun. Solving the equality generalized traveling salesman problem using the Lin–Kernighan–Helsgaun algorithm. *Computer Science Report*, 141, 2014.

[35] K. Helsgaun. Solving the bottleneck traveling salesman problem using the Lin–Kernighan–Helsgaun algorithm. Technical report, Technical Report, Computer Science, Roskilde University, 2014.

[36] W. R. Stewart. Accelerated branch exchange heuristics for symmetric traveling salesman problems. *Networks*, 17(4):423–437, 1987.

[37] P. Schönemann. On two-sided orthogonal Procrustes problems. *Psychometrika*, 33(1):19–33, 1968.

[38] S.W. Hadley, F. Rendl, and H. Wolkowicz. A new lower bound via projection for the quadratic assignment problem. *Mathematics of Operations Research*, 17:727–739, 1992.

[39] K. Anstreicher and H. Wolkowicz. On Lagrangian relaxation of quadratic matrix constraints. *SIAM Journal on Matrix Analysis and Applications*, 22(1):41–55, 2000.

[40] T.H. Cormen. *Introduction to algorithms*. MIT press, 2009.

[41] M. Fiedler. Algebraic connectivity of graphs. *Czechoslovak mathematical journal*, 23(2):298–305, 1973.

[42] M. Fiedler. Laplacian of graphs and algebraic connectivity. *Banach Center Publications*, 25(1):57–70, 1989.

[43] F. Chung. *Spectral graph theory*, volume 92. American Mathematical Soc., 1997.

[44] T. Sahai, A. Speranzon, and A. Banaszuk. Hearing the clusters of a graph: A distributed algorithm. *Automatica*, 48(1):15–24, 2012.

[45] G. Reinelt. TSPLIB–A traveling salesman problem library. *ORSA journal on computing*, 3(4):376–384, 1991.